Synopsis V1.3
Proton Dose and Single Event Effects Testing of the Intel Pentium III (P3) and AMD K7 Microprocessors

Jim Howard[1], Evan Webb[2], Ken LaBel[2], Marty Carts[3], Ron Stattel[3] and Charlie Rogers[3]
1. Jackson and Tull Chartered Engineers, Washington DC 20018
2. NASA GSFC, Greenbelt, MD 20771
3. Raytheon ITSS, Greenbelt, MD 20771

Test Dates: June 21-24, 2000. Report Date: October 27, 2000.

Introduction

As part of the Remote Exploration and Experimentation Project, work was funded for the "Radiation Evaluation of the INTEL Pentium III and Merced Processors and Their Associated Bridge Chips." As a first step in the completion of this work, Intel Pentium III and AMD K7 processors were tested at the proton facility at the Indiana University Cyclotron Facility (IUCF). There were four main objectives for this initial proton testing. They were to do a first order estimate of the total ionizing dose (TID) sensitivity, evaluate the single event latchup (SEL) sensitivity to protons, determine the susceptibility to single-event-induced functional interrupts (SEFI) and the single event upset (SEU) rate for the memory component of the processors. The remainder of this report details the test process, methodology and results from this first look at the P3 technology.

Devices Tested

Pentium III devices with rated clock speeds of 550, 650 and 700 MHz were used. As a last minute test, K7 microprocessors running at 650 and 700 MHz were tested. Pentium III devices were manufactured by Intel and the K7 devices were manufactured by Advanced Micro Devices (AMD). All devices were characterized prior to exposure. A listing of all devices used in this testing are given in the Table below.

Test Facility

Facility: Indiana University Cyclotron Facility
Proton Energy: 189.9 MeV incident on DUT structure
Flux: 2.9 x 10^5 to 5.9 x 10^{10} protons / cm^2 / s.

Test Methods

Temperature:

The test was conducted at room temperature. Intel P3 junction temperature was monitored using an on-die diode.

Test Hardware:

The test system is described from a hardware perspective at three levels; starting with the system level, and followed by the motherboard description and the device under test (DUT) processor description. A description of the functionality of the system concludes this section. These descriptions apply to the Intel Pentium III. The AMD K7 section describes any differences from the P3 setup. Figure 1 illustrates the overall test configuration.

Device Under Test (DUT) Table				
Device	**Vendor**	**Speed**	**DUT Number**	**Package Markings**
Pentium III	Intel	550 MHz	550_2	550/256/100/1.65V S1 90050493-0099 MALAY imc '99 SL3V5
Pentium III	Intel	650 MHz	650_1	650/256/100/1.65V S1 10100418-0176 PHILIPPINES imc '99 SL3KV
Pentium III	Intel	650 MHz	650_2	650/256/100/1.65V S1 10100418-0161 PHILIPPINES imc '99 SL3KV
Pentium III	Intel	650 MHz	650_3	650/256/100/1.65V S1 10100418-0293 PHILIPPINES imc '99 SL3KV
Pentium III	Intel	700 MHz	700_1	700/256/100/1.65V S1 90160187-0108 MALAY imc '99 SL454
Pentium III	Intel	700 MHz	700_2	700/256/100/1.65V S1 90160187-0055 MALAY mc '99 SL454
Pentium III	Intel	700 MHz	700_3	700/256/100/1.65V S1 90160187-0057 MALAY imc '99 SL454
K7	AMD	650 MHz	650_1K	AMD-K7650MTR51B A 230015009833
K7	AMD	700 MHz	700_1K	AMD-K7700MTR51B A 210019614073

The system consists of the motherboard (with power supply, floppy disk drive, monitor and keyboard) and VME Extended Instrumentation (VXI) bus subsystems. A description of the VXI subsystem and the cabling between the subsystems follows.

The VXI subsystem, as shown in Figure 2, consists of the VXI chassis, an embedded controller (running Win98, Labview (LV) environment, and a custom LV application), a signal switch matrix, and two digital multimeters (DMMs) in the voltage measurement mode. The switch matrix provides two functions: The multiplexing of analog signals to one of the DMMs and contact closures (pulling signal levels to ground). The other DMM is dedicated to monitoring one specific analog value.

Most of the cabling between the VXI and the motherboard leaves the VXI from the switch matrix. The exception, a serial (RS-232) cable, brings telemetry from the motherboard directly to the VXI controller. The controller resides in the test chamber. Its user interface (keyboard / monitor / mouse), along with its Ethernet connection, are extended out to the user facility.

The motherboard subsystem consists of a commercial motherboard, a standard PC ATX power supply, a floppy disk drive, RAM module, video card, the DUT processor components and the user interface (keyboard / monitor / mouse). The motherboard resides in the test chamber, positioned so that only the DUT processor is exposed directly to the proton beam. The user interface is extended out to the user facility. Two

motherboard signals, both momentary contact closures, are connected for control by the VXI — motherboard power on/off (MotherPow), and motherboard soft reset (MotherSR).

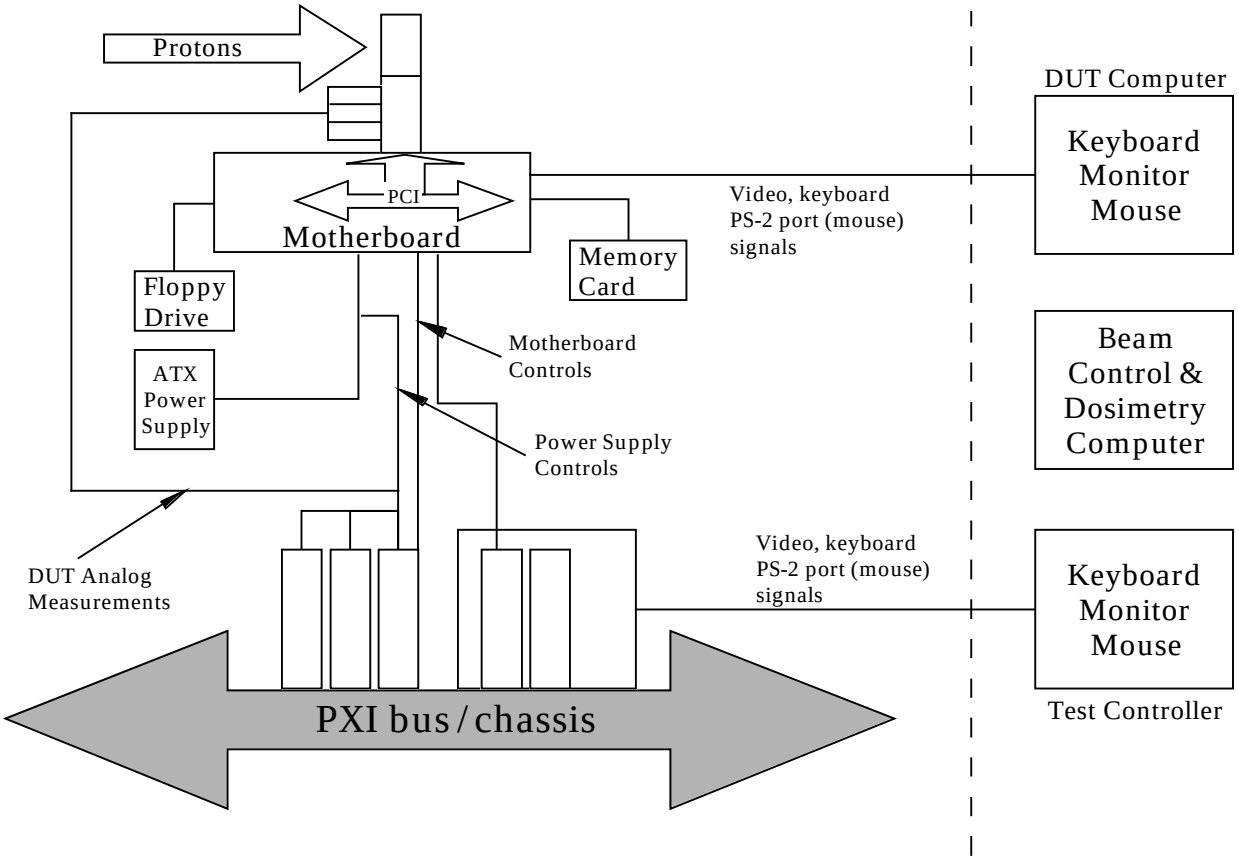

Figure 1. Block diagram illustrating the overall test configuration.

The ATX power supply on/off state is controlled by a signal from the motherboard. This signal (PS_ON#) is, approximately, a latched toggle of MotherPow. MotherPow controls power-down functions on the motherboard, but its latch/toggled version is disconnected from the ATX power supply's PS_ON# so that it can be controlled directly from the VXI.

The DUT software periodically reports to the VXI through the motherboard's serial port via a null modem cable.

Intel P3

The DUT processor, shown in Figure 3, is a modification of a standard Pentium III for the 242-contact slot connector (SC242) module. (The Pentium III is available both in this form and in a 370-pin zero insertion force socket (PGA370) form. The SC242 form is chosen because it is available in higher performance versions and because of packaging/beam-access considerations). The SC242 module consists, in order of arrangement from "back" to front of a plastic backside cover, the printed circuit board (PCB), the heatsink, and a fan/frontside cover. The processor die is mounted facedown onto an organic land grid array (OLGA), a small (1" x 1") PCB (with die mounted directly to it) which is in turn mounted to the PCB with 242 card edge connector pins.

The PCB also carries power distribution traces, high frequency bypass capacitors, processor identification (ID) and voltage identification jumpers. Level 2 cache memory (L2) is incorporated on-die for all P3 processors involved in this test (in some other versions of SC242 module processors L2 resides within plastic packaged ICs mounted to the PCB).

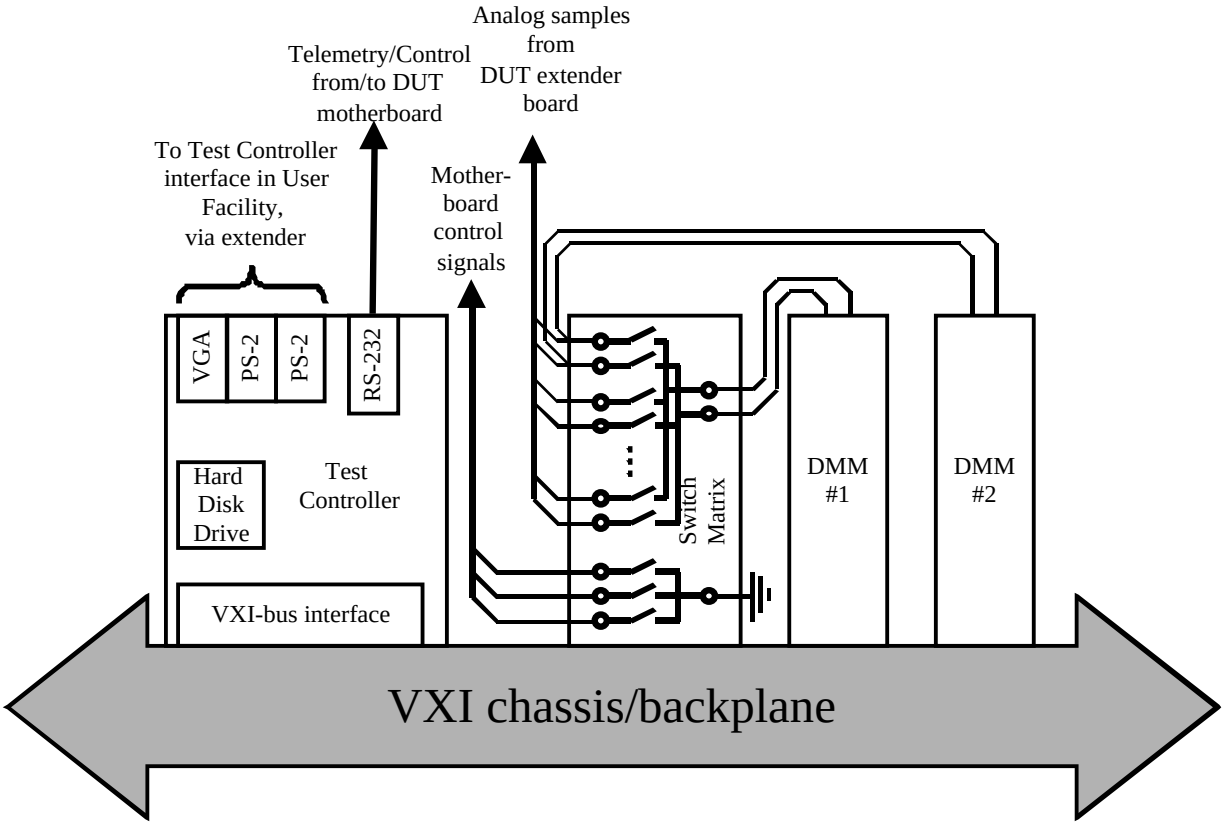

Figure 2. Block diagram of the VXI subsystem.

Figure 3. Photographs of Pentium III SC242 processor showing fan side and front side.

Nominally, the SC242 module DUT processor is inserted (with card-edge fingers downwards) into a motherboard socket. For two reasons, a short (~1" high) extender board was inserted between the DUT processor and the motherboard socket. First, a clear line of site for the proton beam was thus established (connectors and the system

RAM was avoided). Second, the extension provided the best opportunity for monitoring DUT currents. The extender board was modified to insert a low resistance in series with the power traces. Cabling was added to monitor the developed voltages. In addition, the +5Vdc from the socket was brought to the connector, providing for the possibility for amplifiers in the case that the voltages need buffering to overcome high noise levels at the facility. A photograph of the modified extender card can be seen in Figure 4.

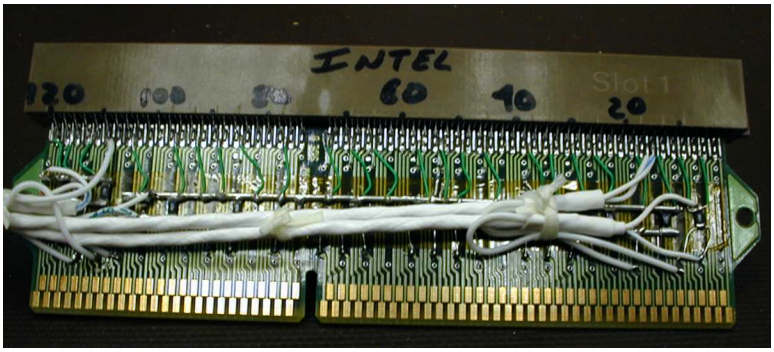

Figure 4. Photograph showing the Pentium III extender card.

Each DUT processor module was disassembled and reassembled with the heatsink/fan assembly translated to a position above the DUT to improve beam access. A thermal transfer plate ("heat pipe") provided thermal conductivity from the DUT to the heatsink/fan assembly. Thermal pad material instead of thermal grease was used to reduce thermal resistance between the thermal plate and the DUT, and the heatsink, in order to avoid having to deal with activated lithium or silicon greases. This structure is shown in Figure 5.

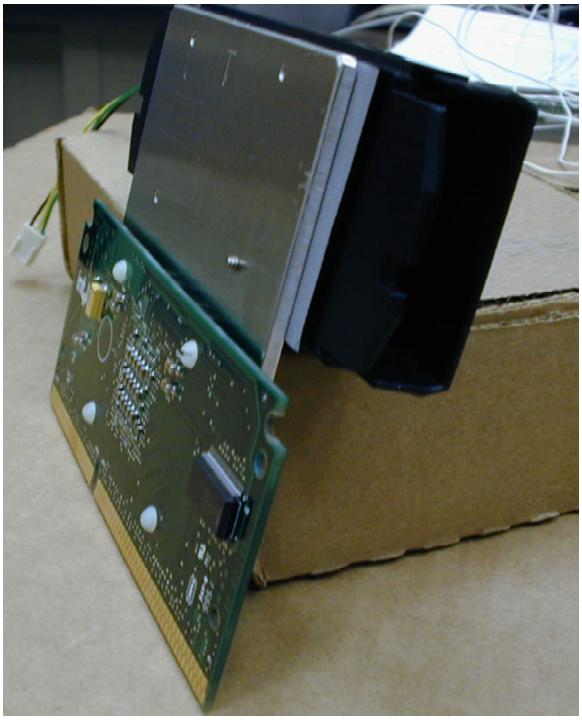

Figure 5. Photograph showing the modified Pentium III processor card.

Signals that are controlled by the VXI subsystem are:

Name	Destination	Description
PS_ON#	ATX Power supply	0V = On; open = Off
MotherPow	Motherboard power switch connector	Pulse low (0 V) to toggle power on and off
MotherReset	Motherboard reset switch connector	Pulse low (0 V) to initiate reset

Signals that are monitored by the VXI are:

Name	Source	Description
I_Vcc_core, Vcc_core	Extender board	Voltage samples of the DUT core current and voltage. Twisted shielded pair (TSP): pair is both sides of the 0.001 ohm current sampling resistor; shield is ground.
I_Vcc_L2, Vcc_L2	Extender board	Voltage samples of the DUT Level 2 cache current and voltage. TSP: pair is both sides of the 0.010 ohm current sampling resistor; shield is ground.
I_Vtt, Vtt	Extender board	Voltage samples of the DUT termination voltage source's current and voltage. TSP: pair is both sides of the 0.010 ohm current sampling resistor; shield is ground.
V_temp	Extender board	Voltage sample of the on-die temperature sensing diode. TSP.
Telemetry	COM1	RS-232 carrying telemetry from the DUT processor.

AMD K7

The DUT processor is a modification of a standard AMD K7. It, as is the Pentium III, is available in the SC242 form but the signal assignments are entirely different, to the extent that the K7 requires the use of a different motherboard, different chipset, and different DUT extender board. The construction of the K7 DUT processor module is similar to that of the Pentium III.

A DUT extender board, which had power group bussing for the Pentium III was modified to un-bus the Pentium III power groups, re-bus them appropriately for the K7, and then to add current sampling capability.

Each DUT processor module was disassembled and reassembled with a thermal plate to allow uniform beam access to the die. This structure is shown in Figure 6.

Signals that are controlled by the VXI subsystem are the same as for the Pentium III (PS_ON#, MotherPow, and MotherReset). However, no signals were monitored by the VXI system, as the motherboard and extender card were modified on-site for the quick test. The AMD does not have an on-die temperature sensing diode.

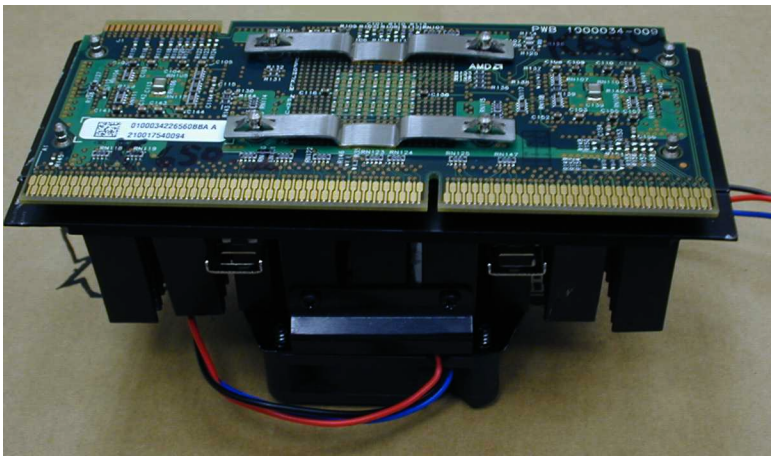

Figure 6. Photograph showing the AMD K7 processor card, including heatsink modifications.

Test Software:

The Pharlap embedded operating system is used by the DUT computer to execute the test code. The test code is written in Microsoft Visual C++ 6.0 professional edition with Pharlap add-ins and Pharlap 386 assembly. Instructions beyond 386 are added with macros. The Pharlap add-ins enable remote debugging of the code through the serial ports.

The software executes with two or more threads. The main thread is executed upon booting the system from the floppy disk. For the task switching test (D), eight other threads are launched to test the switching between threads while the main thread goes into sleep mode checking for test completion every 0.1 seconds. When the test completes the eight iterations, the threads are terminated. The main thread displays a menu and waits for another test to be selected. For all of the other tests, a second thread is launched to run the test software and the main thread goes into sleep as in test D. When the test completes the second thread is terminated. The main thread displays a menu of all available tests to run as outlined below.

The test software sends a keep alive to the VXI and the screen every second. If errors occur, the test software accumulates errors for one second and then dumps error codes to the VXI system and the screen.

Boot process:

- The Pharlap operating system and monitor are loaded from the floppy disk.
- The test software (p3test.exe) is loaded from the floppy disk.
- The test software is executed from random access memory.
- The test software displays all available tests to the VXI and the test computer's monitor.
- The program waits for input from the VXI or the test computer's keyboard.
- The selected test is executed until stopped by an escape from the VXI or the test computer's keyboard.

Description of the six tests currently available in p3test.exe:

A: This test repeats the following steps until stopped. Errors are reported to the VXI and the screen once per second when operating at 660MHz. The instruction timing is used to measure each second.

- All of the registers are initialized to memory at checkData+rCode+20. Where rCode is 0:ebx 4:ecx 8:edx 0xC:ebp 0x10:edi. Currently the memory is set to 0AAAAAAAAh.
- The register r is compared with the memory at checkData+rCode.
- If a miscompare results an error is reported to the error buffer.
- An attempt is made to set the register to the original value and the result of this attempt is stored as the new expected value.
- steps b-d are repeated for each register until one second finishes.
- All errors are reported to the dump RAM, the VXI, and the screen.
- If the escape code is received from the keyboard or the VXI then exit.
- go back to a.

An error consists of the following four dwords:

- rCode to show which register failed: 0:ebx 4:ecx 8:edx 0xC:ebp 0x10:edi
- The contents of the register.
- The expected contents of the register.
- The new contents of the register after attempting to reset the register to its original contents.

B: This test repeats the following tests until stopped. Errors are reported to the VXI and the screen once per second when operating at 660MHz. The instruction timing is used to measure each second.

1. The arguments and expected results are loaded into RAM at checkData.
2. The function f is executed and the result compared to the expected result.
3. If a miscompare results, an error is reported to the error buffer.
4. Steps b and c are repeated for each function until one second finishes.
5. All errors are reported to the dump RAM, the VXI, and the screen.
6. If the escape code is received from the keyboard or the VXI then exit.
7. go back to a.

An error consists of three quadwords:

- The first is the function identifier: 0:fadd 24:fsub 48:fmul 72:fdiv 96:fsqrt
- The second is the result of the operation.
- The third is the expected result.

C: This test loads 100000 locations of memory with an incrementing pattern and then compares the block to the incrementing pattern. It repeats this function until an escape is sent from the keyboard or the VXI. Errors are reported immediately to the VXI and the screen. A keep-alive is sent to the screen and the VXI every second.

D: This test launches seven tasks. Each task loops through incrementing a memory location from 0 to 11 counts. When all of the tasks complete, the memory locations are checked to see if they equal 11. It repeats this function until an escape is sent from the keyboard or the VXI. Errors are reported immediately to the VXI and the screen. A keep-alive is sent to the screen and the VXI every second.

E: This test has 16K of instructions run in order. The instruction sequence is repeatedly incrementing the eax register from 0 to 4; checking after each increment to see that it has

done so. If an error occurs the test is aborted and an error message is reported to the VXI and the screen.

F: This test is similar to test B except that the function tested is as follows:

result=cos(cos(cos(sin(sin(sin(sqrt(sqrt(sqrt(sqrt(sqrt(sqrt(a*b))))))))))))

where a=0.123456789 and b=0.987654321

Notes:

(a) The cache can be turned off or on by pressing '@' from the test menu. The cache state switches between three settings:
- all caches off
- level 1 cache on and level 2 cache off
- both level 1 and level 2 caches on

(b) For Test C, the cache must be turned on.

(c) During the IU test trip, the code was modified to also run on the AMD K7.

(d) During the test trip, a first attempt was made at modifying the code to actively handle exceptions. The last number of runs with the P3 was testing out this code.

Test Methodology

TID Test Process

To do a complete characterization of the P3/K7 (Device Under Test (DUT)) for total dose effects would require numerous parametric measurements that were not planned for this test. This level of characterization is not necessary for the rough approximation required for the follow-on single event testing. To this end, it will be sufficient to monitor all the voltages and currents to the P3 (available through the extender card built for the SEE testing) as well as processor functionality. It should be noted that the currents and voltages were not available for the K7 as the extender card was not able to be built with those measuring points on the short time scale. The AMD K7 was only functionally tested after each dose point.

To understand the issues of operational performance with dose, two of the P3's available (650 MHz) were run through total dose procedure. One will be dosed in the biased state and one in the unbiased state (the powered-on P3 will be monitored for latchup to ensure that latchup is not destroying the device rather than the dose). After it is established which of these conditions is worst case, the highest speed P3 obtained (700 MHz) will be set into that condition and total dose tested. This is done to determine if we can expect differences in total dose response with different design generations. The TID results of the powered-on P3 will be used to set a stopping point for SEE testing for a given DUT as it receives dose.

Single Event Effects Test Process

The SEE test process must include methods to test for all aspects of single event effects (latchup, functional interrupts, upsets, etc.). As a number of these effects are sensitive to the software being run and may be sensitive to numerous other conditions, detailed control of the DUT is required. To this end, an extensive operating system would serve no purpose other than a software overhead that is uncontrollable. Therefore, the boot process is done into a minimal operating system (Pharlap) and a test executive is run that allows testing of the DUT at a very low level. A flow diagram of the main testing process is shown in Figure 7.

The main part of this flow process, after the DUT is placed in a known operating state, is waiting for something to happen and then dealing with it. The flow describing

this process is shown in Figure 8. Three general categories of events are expected: functional interrupts, system resets (radiation induced), and non-fatal errors (some error is produced but it does not immediately induce a functional interrupt or system reset. The remainder of the flow for all of these conditions shows the steps required to gain information about what exactly happened and to recover the DUT to a known state. The “Reboot Testing” called in this routine is simply trying a soft reboot three time. If at that point the system has not recovered, a hard boot is then done.

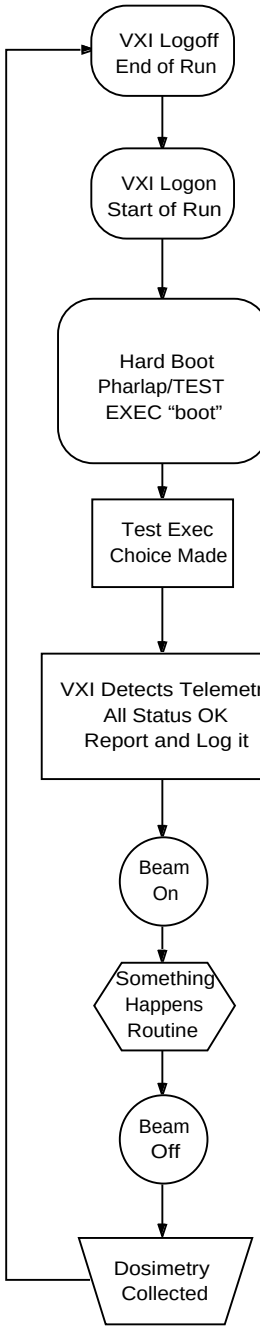

Previous run is completed

Logging is started and a new test run begins.

The test processor is hard-booted (so that a known startup condition exists), the minimal OS and test software is started.

Choice is made as to which test routine is to be run.

Test hardware detects that P3 is up and running test software correctly. This is logged and operator is informed that test is ready to go.

Proton Beam is turned on.

Software loop executes while we wait for something to happen. This is then handled in the “Something Happened Routine”.

When the run is completed, the beam is turned off.

Beam dosimetry is collected and logged with the test results.

Figure 7. Flow diagram and description for main testing loop.

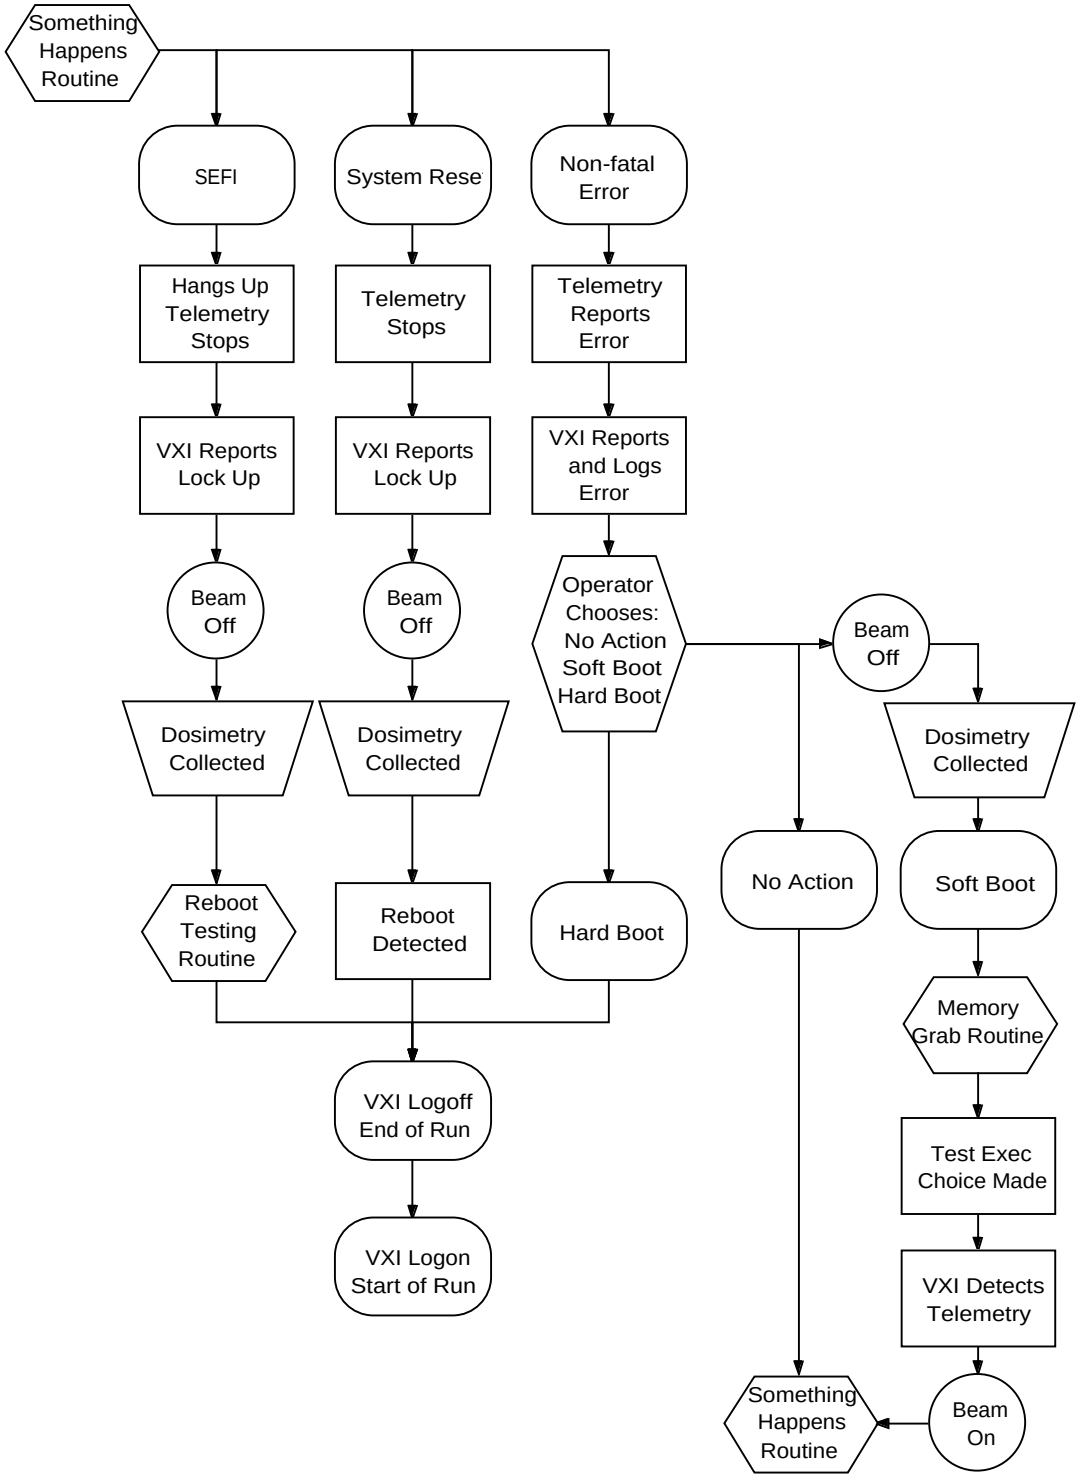

Figure 8. Flow process after an event has been detected.

Results

Total Dose

Three Intel P3 parts were exposed to the total dose environment. The first part was a 650 MHz P3 and it was exposed to protons in an unbiased state (it was placed in the motherboard socket but no power was supplied to the motherboard). It was exposed to 25 krad (Si) of protons in 5 krad (Si) increments. After each dose point, the P3 passed all functional tests and the monitored voltages and currents did not change. An example of the monitored voltages and currents is shown in Figure 9.

The second part to be exposed was a 650 MHz P3 that was biased and in nominal operation. It was exposed to 50 krad (Si) of protons in 9 increments. After each dose point, the P3 passed all functional tests and the monitored voltages and currents did not change. Functional testing consisted of booting the processor and executing every option in the program executive code.

The final P3 part tested was a biased, operating 700 MHz processor. It was exposed to 100 krad (Si) of protons in 20 krad (Si) increments. After each dose point, the P3 passed all functional tests and the monitored voltages and currents did not change.

One AMD K7 part was used for a total dose determination. The 650 MHz K7 was exposed to protons in an unbiased state (it was just clamped in place in the beam line). It was exposed to 100 krad (Si) of protons in 8 increments. After each dose point, the K7 passed all functional tests.

This testing seems to indicate that this generation of P3 and K7 processors are TID hard to greater than 100 krad (Si). It should be pointed out that no effort was made at this point to do parametric timing measurements. It is these that are expected to be the most sensitive to dose. The parts, though, obviously did not degrade in timing sufficiently to fail any of the functional tests that were performed.

Single Event Latchup

Five different P3 processors (one 550 MHz, two 650 MHz and two 700 MHz) run at eight different clock speeds (one each of the 550, 650 and 700 MHz were clocked down to 366, 433 and 466 MHz, respectively). During these tests the processors were running one of the tests in the test executive (tests were varied) and exposed to proton fluences (per run) that varied from 2.7×10^5 to 3.3×10^{11} protons/cm^2. The P3 parts were tested in 107 different conditions (processor speed, cache on/off, and software executing) and the K7 parts were tested in 7 different conditions. In all of the testing, no evidence of latchup was observed (presence of latchup would be indicated by a sharp increase in I_Vcc_core). In fact, no statistically significant increase in that or any current was observed during any of the testing. Nor were there any cases that required a hard boot to recover the system to normal operation.

Single Event Upsets – Non-SEFI

Register and cache test, as described in the software setup section, were run to determine the upset sensitivity of these components. Direct evidence of these events was observed through the execution of the software (leading to non-fatal errors). In addition, some of the observed Single Event Functional Interrupts (SEFIs) were a direct result of upsets in the registers and/or caches. The software used in this testing was set up to monitor for these events but not well set up to deal with the SEFIs, to be described next. As these SEFI events would interrupt the SEU monitoring at random times the

collection of SEU rates was limited. The software has been improved and the SEU data collection should be better in the next set of proton tests.

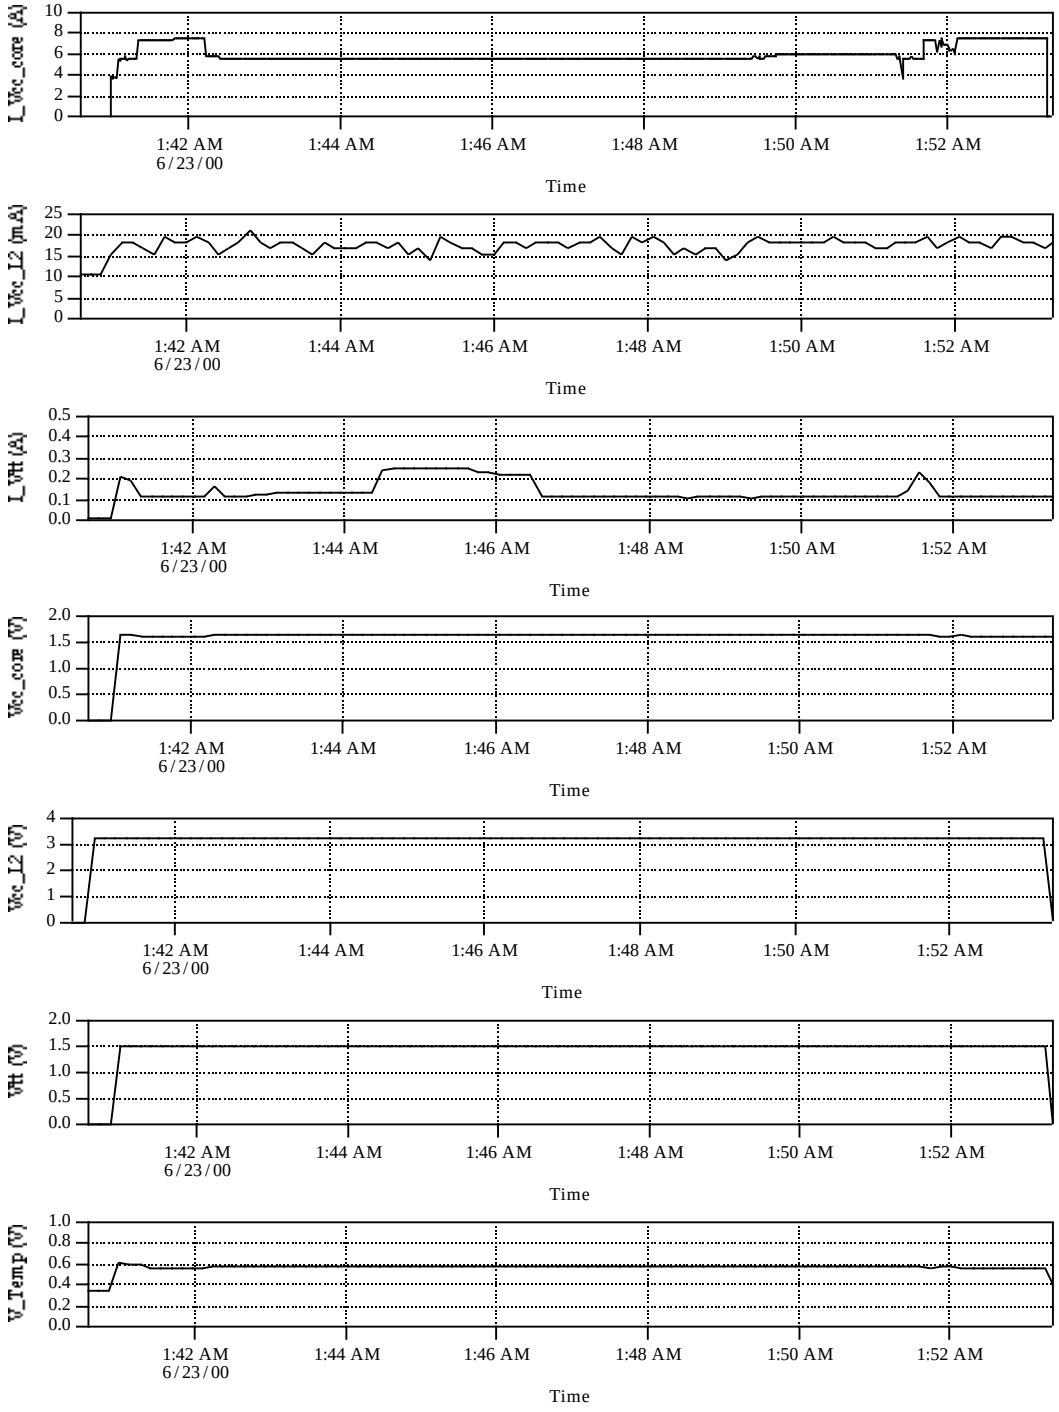

Run #026. Intel P3 700 MHz. DUT #700_3

Figure 9. Example Strip Chart for the monitored voltages and currents.

Single Event Functional Interrupts (SEFI)

During the testing mentioned in the above latchup section, approximately 100 of those tests were classified as single event upset testing. In almost every one of these 100 cases, a SEFI was observed. In all of these cases, the proton beam was stopped as soon as there was an indication that a SEFI had occurred (exception returned, telemetry stopped, etc.). The cross section for an event is normally defined as

$$\sigma = \frac{N}{F},$$

where σ is the cross section for the event, N is the number of events for the test fluence, and F is the fluence for that test run.

Based on this, a cross section can be defined to be one over the fluence required to cause the SEFI (it is one over the proton fluence, as there was only one SEFI event). It should be noted that what was termed as a SEFI was when the executive program that was running, terminated with an exception or failed in sending out the proper telemetry, indicating that it was still alive. The issue that arose during the testing was whether the processor was truly hung after an exception was received. The software that was written for this testing did not have exception handling. Therefore, when an exception was generated, the operating system took over and dumped the program. It is possible that the processor was still functioning properly and, if the code would have been restarted, it would have continued normally. An attempt was made during the testing to rewrite the test executive to include exception handling. While this first version did have its problems, it did indicate that at least some of the exceptions were not fatal and the test program could be continued. This is a very important point as it could have significant meaning for software-level mitigation.

All the cases that led to a SEFI event were placed into categories so that the data could be looked at from different variables. The categories are related to the devices tested and the status of the cache during the test. It should be noted that the data was categorized by the test software run, but no significant difference was noted there. The average and standard deviation of the data are presented in the following table.

The first column of this table contains the average value for that category. The second contains the standard deviation (a zero in this column indicates that only one data point was taken for that category). The third column contains information about the device being tested. The one entry with K7 indicates that device is an AMD K7 microprocessor. All other entries are for Intel Pentium III devices. “All Devices” indicates that the average was taken over all devices independent of the DUT rated speed. If a device speed is indicated, then that category is for a device operating at that speed only (366 MHz is a clocked-down 550 MHz, 433 MHz is a clocked-down 650 MHz, and 466 MHz is a clocked-down 700 MHz). The fourth column contains the status of the cache during the testing. “Cache On” indicates that both the L1 and on-die L2 cache were in use and “Cache Off” indicates that neither the L1 nor on-die L2 cache was in use. “L1 Cache” indicates that only the L1 cache was in use, and “All SEFIs” indicating an average of all events independent of the cache state). While a table is useful for see the numbers, a graph often better shows what the data is saying. Figure 10 the data from the above table plotted as a function of the cache state, where each device category is plotted as a different symbol.

SEFI Cross Section Results				
Average (cm^2)	**Standard Deviation**	**Operating Speed**	**Case**	**Device Under Test Speed**
3.43E-09	5.41E-09	366, 550 MHz	All SEFIs	550 MHz
3.50E-09	5.56E-09	433, 650 MHz	All SEFIs	650 MHz
7.58E-10	7.43E-10	466, 700 MHz	All SEFIs	700 MHz
2.71E-09	4.74E-09	All Speeds	All SEFIs	550, 650, 700 MHz
1.83E-10	5.04E-11	366 MHz	Cache Off	550 MHz
5.76E-10	2.10E-10	433 MHz	Cache Off	650 MHz
1.03E-10	7.93E-11	466 MHz	Cache Off	700 MHz
2.68E-10	1.70E-10	550 MHz	Cache Off	550 MHz
4.41E-10	2.95E-10	650 MHz	Cache Off	650 MHz
2.51E-10	3.64E-10	700 MHz	Cache Off	700 MHz
3.24E-10	2.97E-10	All Speeds	Cache Off	550, 650, 700 MHz
4.89E-09	6.28E-09	366 MHz	Cache On	550 MHz
6.41E-09	8.95E-09	433 MHz	Cache On	650 MHz
9.72E-10	3.19E-10	466 MHz	Cache On	700 MHz
6.05E-09	6.44E-09	550 MHz	Cache On	550 MHz
4.94E-09	6.99E-09	650 MHz	Cache On	650 MHz
1.35E-09	6.20E-10	700 MHz	Cache On	700 MHz
4.41E-09	5.94E-09	All Speeds	Cache On	550, 650, 700 MHz
3.00E-09	3.50E-09	K7 700 MHz	Cache On	K7 700 MHz
5.92E-09	0.00E+00	433 MHz	L1 Cache	650 MHz
1.53E-09	0.00E+00	550 MHz	L1 Cache	550 MHz
5.52E-09	5.58E-10	650 MHz	L1 Cache	650 MHz
3.19E-09	2.77E-09	All Speeds	L1 Cache	550, 650, 700 MHz

It should be noted here, as well as for the next set of figures, that error bars (from the standard deviation column) are not included. These were excluded because, due to their length, confused reading of the symbols. Also, because of their length, solid statistically significant conclusions are not possible. However, while it is not possible to say that the "Cache Off" category is an order of magnitude better that any of the other cases, it is the case where the P3 is least likely to have a SEFI event. What this data also shows is that there is little difference between the "Cache Off" and the "L1 Cache" cases. This indicates that if there is a real difference between the cache on and off cases, it lies in the L2 cache mainly.

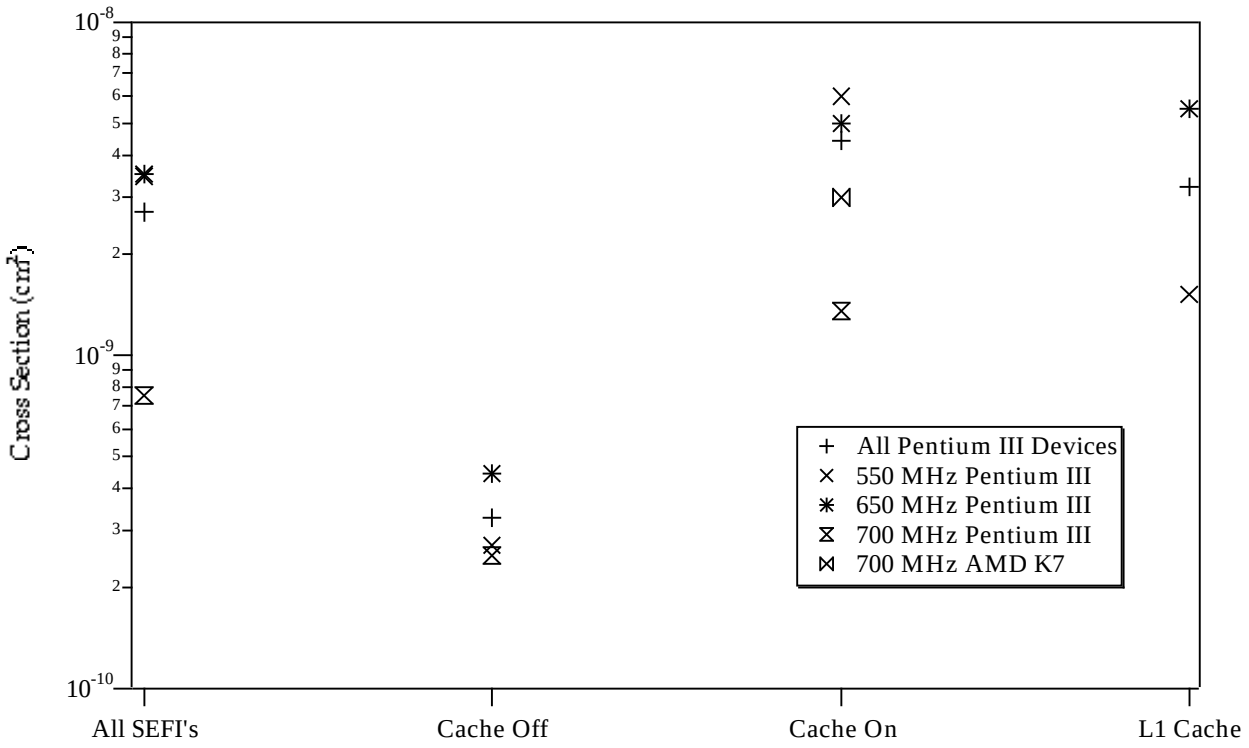

Figure 10. Plot showing the SEFI cross section as a function of the cache state during testing.

To look at the two cache cases more closely, Figure 11 gives that “Cache On” case as a function of the device type and Figure 12 gives that “Cache On” case as a function of the device type. As with the previous figure, the error bars on this plot can span the entire height of the graph for some data points. With this in mind, it is still interesting to note that the two data points that are the lowest in Figure 11 are the 466 MHz and the 700 MHz. It should be remembered that the 466 MHz device is just the clocked-down 700 MHz device. This is not the case for Figure 12 where they are more or less in the same cross section regime as the other devices. This could be an indication that the on-die L2 cache used in the 700 MHz device is different than that used in the 550 or 650 MHz. Much better statistics would be needed to make that statement with any confidence, though.

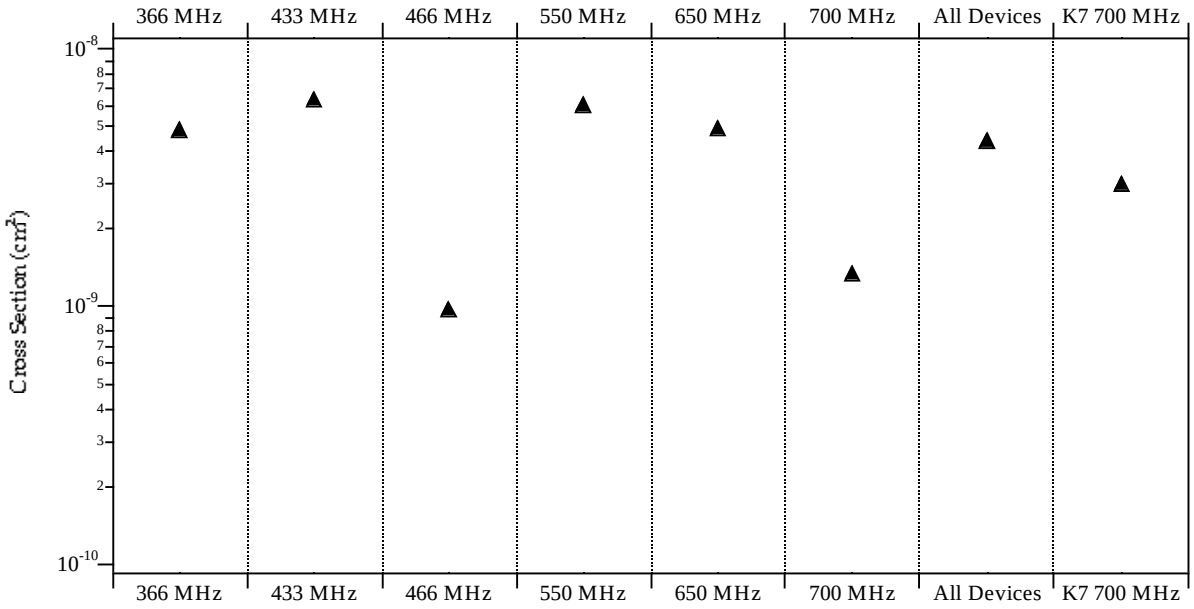

Figure 11. Plot showing the SEFI cross section as a function the device category for the case of “Cache On”.

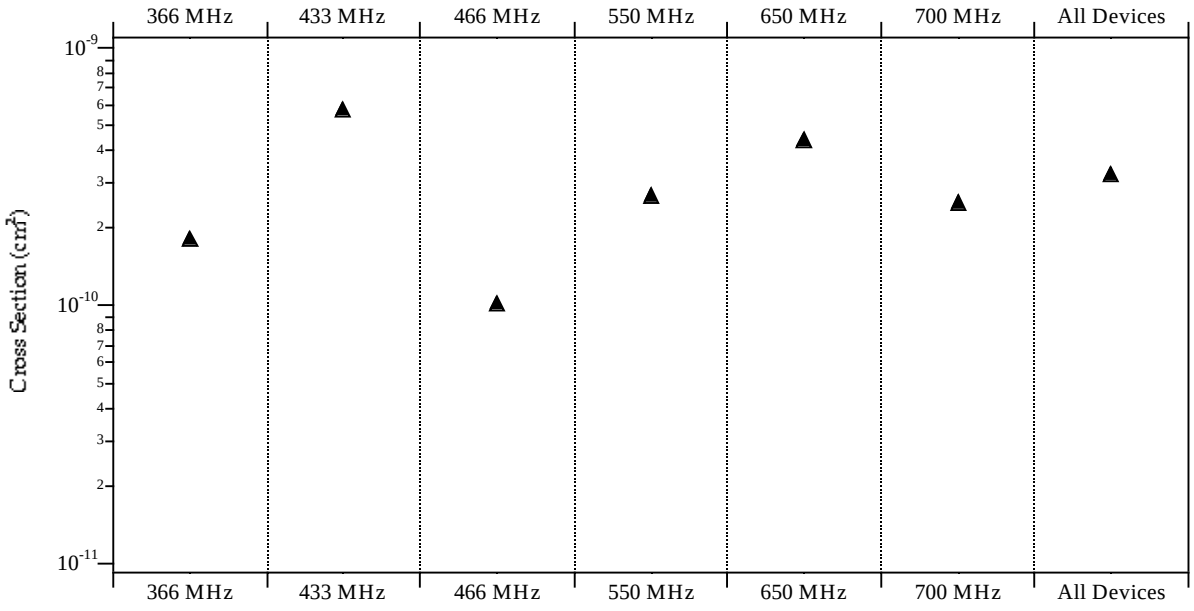

Figure 12. Plot showing the SEFI cross section as a function the device category for the case of "Cache Off".

As a final discussion on the SEFI testing, some testing was done doing a reboot while exposed to the proton beam. One of the SEFI events during a test was an event that triggered a system reboot. The system was allowed to reboot without turning the

proton beam off. However, the system did not fully reboot before another SEFI occurred that halted the reboot process. A series of proton flux rate tests were then run where the proton flux was lowered, five reboot test were done under proton irradiation, and the number of successful reboots were recorded. This was done again for a slightly higher proton flux until there were zero-out-of-five successful reboots. It was found that at a flux of 3×10^6 protons/cm^2-sec, reboots could be done consistently without any failures. However, by a flux of 5×10^6 protons/cm^2-sec failures were beginning to start and by 3×10^7 protons/cm^2-sec there were no successful reboots in five attempts. It should be noted that most of the SEE runs were done at a flux rate of approximately 1×10^8 protons/cm^2-sec. These levels of proton flux would not be expected in the nominal space environment, but could be approached during a large solar particle event.